

Particle algorithms for optimization on binary spaces

Christian Schäfer^{1,2}

April 9, 2012

We discuss a unified approach to stochastic optimization of pseudo-Boolean objective functions based on particle methods, including the cross-entropy method and simulated annealing as special cases. We point out the need for auxiliary sampling distributions, that is parametric families on binary spaces, which are able to reproduce complex dependency structures, and illustrate their usefulness in our numerical experiments. We provide numerical evidence that particle-driven optimization algorithms based on parametric families yield superior results on strongly multi-modal optimization problems while local search heuristics outperform them on easier problems.

Keywords Binary parametric families · Unconstrained binary optimization · Sequential Monte Carlo · Cross-Entropy method · Simulated Annealing

1 Particle optimization

1.1 Introduction

1.1.1 Pseudo-Boolean optimization

We call $f: \mathbb{B}^d := \{0, 1\}^d \rightarrow \mathbb{R}$ a pseudo-Boolean function. The present work discusses approaches to obtain heuristics for the program

$$\begin{aligned} & \text{maximize} && f(\mathbf{x}) \\ & \text{subject to} && \mathbf{x} \in \mathbb{B}^d \end{aligned} \tag{1}$$

using sequential Monte Carlo techniques. In the sequel, we refer to f as the *objective function*. For an excellent overview of applications of binary programming and equivalent problems we refer to the survey paper by [Boros and Hammer \(2002\)](#) and references therein.

The idea to use particle filters for global optimization is not new [([Del Moral et al., 2006](#)), Section 2.3.1.c], but novel sequential Monte Carlo methodology making use of suitable parametric families on binary spaces [Schäfer and Chopin \(2011\)](#) may allow to construct more efficient samplers for the special case of pseudo-Boolean optimization.

¹Centre de Recherche en Économie et Statistique, 3 Avenue Pierre Larousse, 92240 Malakoff, France

²Centre de REcherches en MATHématiques de la DEcision, Université Paris-Dauphine, Place du Maréchal de Lattre de Tassigny 75775 Paris, France

We particularly discuss how this methodology connects with the cross-entropy method [Rubinstein \(1997\)](#) which is another particle driven optimization algorithm based on parametric families.

The sequential Monte Carlo algorithm as developed by [Schäfer and Chopin \(2011\)](#) is rather complex compared to local search algorithms such as simulated annealing [Kirkpatrick et al. \(1983\)](#) or k -opt local search [Merz and Freisleben \(2002\)](#) which can be implemented in a few lines. The aim of this paper is to motivate the use of advanced particle methods and sophisticated parametric families in the context of pseudo-Boolean optimization and to provide conclusive numerical evidence that these complicated algorithms can indeed outperform simple heuristics if the objective function has poorly connected strong local maxima. This is not at all clear, since, in terms of computational time, multiple randomized restarts of fast local search heuristics might very well be more efficient than comparatively complex particle approaches.

1.1.2 Outline

The article is structured as follows. We first introduce some notation and review how to model an optimization problem (1) as a filtering problem on an auxiliary sequence of probability distributions. Section 2 describes a sequential Monte Carlo sampler [Del Moral et al. \(2006\)](#) designed for global optimization on binary spaces [Schäfer and Chopin \(2011\)](#). Section 3 reviews three parametric families for sampling multivariate binary data which can be incorporated in the proposed class of particle algorithms. Section 4 discusses how the cross-entropy method [Rubinstein \(1997\)](#) and simulated annealing [Kirkpatrick et al. \(1983\)](#) can be interpreted as special cases of the sequential Monte Carlo sampler. In Section 5 we carry out numerical experiments on instances of the unconstrained quadratic binary optimization problem. First, we investigate the performance of the proposed parametric families in particle-driven optimization algorithms. Secondly, we compare variants of the sequential Monte Carlo algorithm, the cross-entropy method, simulated annealing and simple multiple-restart local search to analyze their respective efficiency in the presence or absence of strong local maxima.

1.1.3 Notation

We briefly introduce some notation that might be non-standard. We denote scalars in italic type, vectors in italic bold type and matrices in straight bold type. Given a set M , we write $|M|$ for the number of its elements and $\mathbb{1}_M$ for its indicator function. For $a, b \in \mathbb{Z}$ we denote by $\llbracket a, b \rrbracket = \{a, \dots, b\}$ the discrete interval from a to b . Given a vector $\mathbf{x} \in \mathbb{B}^d$ and an index set $I \subseteq \llbracket 1, d \rrbracket$, we write $\mathbf{x}_I \in \mathbb{B}^{|I|}$ for the sub-vector indexed by I and $\mathbf{x}_{-I} \in \mathbb{B}^{d-|I|}$ for its complement. We occasionally use the norms $\|\mathbf{x}\|_\infty := \max_i x_i$ and $|\mathbf{x}| := \sum_{i=1}^d |x_i|$.

1.2 Statistical modeling

1.2.1 Associated probability measures

For particle optimization, the common approach is defining a family of probability measures $\{\pi_\varrho : \varrho \geq 0\}$ associated to the optimization problem $\max_{\mathbf{x} \in \mathbb{B}^d} f(\mathbf{x})$ in the sense

that

$$\pi_0 = \mathcal{U}_{\mathbb{B}^d}, \quad \lim_{\varrho \rightarrow \infty} \pi_\varrho = \mathcal{U}_{M_f},$$

where \mathcal{U}_S denotes the uniform distribution on the set S and $M_f = \operatorname{argmax}_{\mathbf{x} \in \mathbb{B}^d} f(\mathbf{x})$ the set of maximizers. The idea behind this approach is to first sample from a simple distribution, potentially learn about the characteristics of the associated family and smoothly move towards distributions with more mass concentrated in the maxima. We review two well-known techniques to explicitly construct such a family π_ϱ .

Tempered family We call $\{\pi_\varrho: \varrho \geq 0\}$ a tempered family, if it has probability mass functions of the form

$$\pi_\varrho(\gamma) := \nu_\varrho \exp(\varrho f(\gamma)), \quad (2)$$

where $\nu_\varrho^{-1} := \sum_{\gamma \in \mathbb{B}^d} \exp(\varrho f(\gamma))$.

As ϱ increases, the modes of π_ϱ become more accentuated until, in the limit, all mass is concentrated on the set of maximizers. The name reflects the physical interpretation of $\pi_\varrho(\mathbf{x})$ as the probability of a configuration $\mathbf{x} \in \mathbb{B}^d$ for an inverse temperature ϱ and energy function $-f$. This is the sequence used in simulated annealing [Kirkpatrick et al. \(1983\)](#).


Level set family We call $\{\pi_\varrho: \varrho \geq 0\}$ a level set family, if it has probability mass functions of the form

$$\pi_\varrho(\gamma) := |L_\varrho^+|^{-1} \mathbb{1}_{L_\varrho^+}(\gamma), \quad (3)$$

where $L_\varrho^+ := \{\gamma \in \mathbb{B}^d: \varrho[f(\mathbf{x}^*) - f(\gamma)] \leq 1\}$ for $\mathbf{x}^* \in M_f$.

Indeed, L_ϱ^+ is the super-level set of f with respect to the level $c = f(\mathbf{x}^*) - 1/\varrho$, for $\varrho > 0$, and $\pi_\varrho(\gamma)$ is the uniform distribution on L_ϱ^+ . As ϱ increases, the support of π_ϱ becomes restricted to the points that have an objective value sufficiently close to the maximum of the f . In the limit, the support is reduced to the set of global maximizers.

Figure 1: Associated sequences π_{ϱ_t} for a toy example $f: \mathbb{B}^3 \rightarrow [-20, 20]$. The colors indicate the advance of the sequences from yellow to red. For simplicity, we choose $\varrho_t = t$ for $t \in \llbracket 0, 16 \rrbracket$.



The particle-driven optimization algorithms are computationally more involved than local search heuristics since we need to construct a sequence of distributions instead

of a sequence of states. We shall see that this effort pays off in strongly multi-modal scenarios, where even sophisticated local search heuristics can get trapped in a subset of the state space.

1.2.2 Rare event simulation

While the tempered sequence is based on a physical intuition, the level set sequence has an immediate interpretation as a sequence of rare events since, as ϱ increases, the super-level set becomes a ‘rare event’ with respect to the uniform measure. Rare event simulation and global optimization are therefore closely related concepts and methods for rare event estimation can often be adapted to serve as optimization algorithms.

Particle algorithms for rare event simulation include the cross-entropy method [Rubinstein \(1997\)](#) and the sequential Monte Carlo sampler [Johansen et al. \(2006\)](#). The former uses the level set sequence, the latter uses a *logistic potential family*

$$\pi_\varrho(\gamma) := \nu_\varrho \ell(\varrho[f(\gamma) - f(\mathbf{x}^*)]),$$

where $\nu_\varrho^{-1} := \sum_{\gamma \in \mathbb{B}^d} \ell(\varrho[f(\gamma) - f(\mathbf{x}^*)])$ and $\ell: \mathbb{R} \rightarrow (0, 1)$, $\ell(x) = [1 + \exp(-x)]^{-1}$ denotes the logistic function. [Johansen et al. \(2006\)](#) did not specifically design their algorithm for optimization but their approach to static rare event simulation is closely related to the particle optimization framework.

2 Sequential Monte Carlo

We discuss a static sequential Monte Carlo sampler that uses a transition kernel with independent proposals in the move step based on suitable parametric families. This methodology has been demonstrated to reliably estimate the mean of the posterior distribution in Bayesian variable selection problems for linear normal models [Schäfer and Chopin \(2011\)](#). In this section, we provide a self-contained description of this framework. For a more general overview of sequential Monte Carlo methods we refer to [Del Moral et al. \(2006\)](#).

2.1 Sequential Importance Sampling

For convenience of notation, we index the sequence of distributions $(\pi_{\varrho_t})_t$ directly by t rather than by the parameter of the family ϱ_t . We refer to $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{B}^{n \times d}$ and $\mathbf{w} \in [0, 1]^n$ with $|\mathbf{w}| = 1$ as a *particle system* with n particles. We say the particle system (\mathbf{w}, \mathbf{X}) *targets* the probability distribution π if the empirical distribution $\sum_{k=1}^n w_k \delta_{\mathbf{x}_k}$ converges to π for $n \rightarrow \infty$.

2.1.1 Importance weights

We sample $(\mathbf{w}_0, \mathbf{X}_0)$ with $\mathbf{x}_{1,0}, \dots, \mathbf{x}_{n,0} \sim \mathcal{U}_{\mathbb{B}^d} = \pi_0$ and set $\mathbf{w}_0 = \mathbf{1}/n$ to initialize the system. Suppose we are given a particle approximation $(\mathbf{w}_t, \mathbf{X}_t)$ of π_t and want to target the subsequent distribution π_{t+1} . For all $k \in \llbracket 1, n \rrbracket$ and $\alpha > 0$, we update the weights to

$$w_{k,t,\alpha} := \frac{u_{k,t,\alpha}}{\sum_{i=1}^n u_{i,t,\alpha}}, \text{ where } u_{k,t,\alpha} := w_{k,t} \frac{\tilde{\pi}_{\varrho_t+\alpha}(\mathbf{x}_{k,t})}{\tilde{\pi}_{\varrho_t}(\mathbf{x}_{k,t})}, \quad (4)$$

and $\tilde{\pi}_\varrho \propto \pi_\varrho$ denotes the unnormalized version of π_ϱ . The normalizing constants ν_ϱ and $|L_\varrho^+|$ defined in equations (2) and (3) are unknown but the algorithm only requires ratios

of unnormalized probability mass functions. We refer to α as the *step length* at time t . After updating, the particle system targets the distribution

$$\pi_{\varrho_t+\alpha} \approx \sum_{k=1}^n w_{k,t,\alpha} \delta_{\mathbf{x}_{t,k}}. \quad (5)$$

As we choose α larger, that is $\pi_{\varrho_t+\alpha}$ further from π_{ϱ_t} , the weights become more uneven and the accuracy of the importance approximation deteriorates. If we repeat the weighting step, we just increase α and finally obtain an importance sampling estimate of $\pi_\infty = \mathcal{U}_{M_f}$ with instrumental distribution $\pi_0 = \mathcal{U}_{\mathbb{B}^d}$. This yields a poor estimator since the probability to hit the set M_f with n uniform draws is $1 - (1 - 2^{-d}|M_f|)^n$ and decreases rapidly as the dimension d grows. The pivotal idea behind sequential Monte Carlo is to alternate moderate updates of the importance weights and improvements of the particle system via resampling and Markov transitions.

Effective sample size The importance weight degeneracy is often measured through the *effective sample size* criterion [Kong et al. \(1994\)](#) defined as

$$\eta_n(\mathbf{w}) := [n \sum_{k=1}^n w_k^2]^{-1} \in [1/n, 1].$$

The effective sample size is 1 if the weights are uniform, that is equal to $1/n$; the effective sample size is $1/n$ if all mass is concentrated in a single particle.

2.1.2 Finding the step length

Given any increasing sequence $(\pi_{\varrho_t})_t$, we could repeatedly reweight and monitor whether the effective sample size falls below a critical threshold. For the special case of annealing via sequential Monte Carlo, however, the effective sample size after weighting $\eta_n(\mathbf{w}_{t,\alpha})$ is merely a function of α . For a particle system $(\mathbf{X}_t, \mathbf{w}_t)$ at time t , we pick a step length α such that

$$\eta_n(\mathbf{w}_t) \beta = \eta_n(\mathbf{w}_{t,\alpha}), \quad (6)$$

that is we lower the effective sample with respect to the current particle approximation by some fixed ratio $\beta \in (0, 1)$ [[Jasra et al., 2011](#)], [[Del Moral et al., 2011](#)]. This ensures a ‘smooth’ transition between two auxiliary distributions, in the sense that consecutive distributions are close enough to approximate each other reasonably well using importance weights; in our numerical experiments, we took $\beta = 9/10$. We obtain the associated sequence $(\varrho_t)_t$ by setting $\varrho_{t+1} = \varrho_t + \alpha_t$ where α_t is a unique solution of (6). Since $\eta_n(\mathbf{w}_{t,\alpha})$ is continuous and monotonously decreasing in α we can use bi-sectional search to solve (6).

2.2 Conditioning the particle system

In this section, we discuss how to condition the weighted system and improve its quality before we proceed with the next weighting step.

2.2.1 Resampling

We replace the system $(\mathbf{w}_{t+1}, \mathbf{X}_t)$ targeting π_{t+1} by a selection of particles $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n$ drawn from the current particle reservoir $\mathbf{x}_{1,t}, \dots, \mathbf{x}_{n,t}$ such that

$$\mathbb{E}(n(\mathbf{x}_k)) = n w_k,$$

where $n(\mathbf{x})$ denotes the number of particles identical with \mathbf{x} . Thus, in the resampled system particles with small weights have vanished while particles with large weights have been multiplied. For the implementation of the resampling step, there exist several recipes. We could apply a multinomial resampling [Gordon et al. \(1993\)](#) which is straightforward. There are, however, more efficient ways like residual [Liu and Chen \(1998\)](#), stratified [Kitagawa \(1996\)](#) and systematic resampling [Carpenter et al. \(1999\)](#). We use the latest in our simulations, see Procedure 1.

Procedure 1: Resampling (systematic)

```
[ht] Input:  $\mathbf{w} = (w_1, \dots, w_n)$ ,  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ 
 $v \leftarrow n \mathbf{w}$ ,  $i \leftarrow 1$ ,  $c \leftarrow v_1$ 
sample  $u \sim \mathcal{U}_{[0,1]}$ 
for  $k \in \llbracket 1, n \rrbracket$  do
  | while  $c < u$  do  $i \leftarrow i + 1$ ,  $c \leftarrow c + v_i$ 
  |  $\hat{\mathbf{x}}_k \leftarrow \mathbf{x}_i$ ,  $u \leftarrow u + 1$ 
end
return  $\hat{\mathbf{X}} = (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n)^\top$ 
```

2.2.2 Moving the system

If we repeated the weighting and resampling steps several times, we would rapidly reduce the number of different particles to a very few. The key to fighting the depletion of the particle reservoir is moving the particles according to a Markov transition kernel κ_{t+1} with invariant measure π_{t+1} . The particle $\hat{\mathbf{x}}_{k,t+1}^{(0)}$ is by construction approximately distributed according to π_{t+1} , and a draw

$$\hat{\mathbf{x}}_{k,t+1}^{(1)} \sim \kappa_{t+1}(\bullet \mid \hat{\mathbf{x}}_{k,t+1}^{(0)})$$

is therefore again approximately distributed according to π_{t+1} . The last sample of the generated Markov chain $(\hat{\mathbf{x}}_{k,t+1}^{(0)}, \dots, \hat{\mathbf{x}}_{k,t+1}^{(s)})$ is, for sufficiently many move steps $s \in \mathbb{N}$, almost exactly distributed according to the invariant measure π_{t+1} and independent of its starting point.

2.2.3 Stopping rule

While we could always apply a fixed number of move steps, we rather use an adaptive stopping criterion based on the number of distinct particles.

Particle diversity We define the *particle diversity* as

$$\zeta_n(\mathbf{X}) := n^{-1} |\{\mathbf{x}_k : k \in \llbracket 1, n \rrbracket\}| \in [1/n, 1].$$

Ideally, the sample diversity $\zeta_n(\mathbf{X})$ should correspond to the expected diversity

$$\zeta_n(\pi) := 1 \wedge n^{-1} \sum_{\gamma \in \mathbb{B}^d} \mathbb{1}_{\{\mathbf{x} \in \mathbb{B}^d : c_n \pi(\mathbf{x}) \geq 1\}}(\gamma),$$

where c_n is the smallest value that solves $\sum_{\gamma \in \mathbb{B}^d} \lfloor c_n \pi(\gamma) \rfloor \geq n$. This is the particle diversity we would expect if we had an independent sample from $\pi(\mathbf{x})$. Therefore, if κ_{t+1} is fast-mixing, we want to move the system until

$$\zeta_n(\hat{\mathbf{X}}_{t+1}^{(s)}) \approx \zeta_n(\pi_{t+1}).$$

Since the quantity on the right hand side is unknown, we stop moving the system as soon as the particle diversity reaches a steady state we cannot push it beyond [Schäfer and Chopin \(2011\)](#).

More precisely, we stop if the absolute diversity is above a certain threshold $\zeta^* \approx 0.95$ or the last improvement of the diversity is below a certain threshold $\zeta_\Delta^* > 0$. We always stop after a finite number of steps but the thresholds ζ^* and ζ_Δ^* need to be calibrated to the efficiency of the transition kernel. For slow-mixing kernels, we recommend to perform batches of consecutive move steps instead of single move steps.

If the average acceptance rate $\bar{\lambda}$ of the kernel (see Section 2.2.4) is smaller than ζ_Δ^* , it is likely that the algorithm stops after the first iteration although further moves would have been necessary. We could adaptively adjust the threshold ζ_Δ^* to be proportional to an estimate of the average acceptance rate; for our numerical experiments, however, we kept it fixed to $\zeta_\Delta^* \approx 10^{-2}$.

Procedure 2: Move

Input: $\mathbf{X} = (\mathbf{x}_1^{(0)}, \dots, \mathbf{x}_n^{(0)})^\top$ **targeting** π
 $\kappa(\gamma \mid \bullet)$ with $\pi(\gamma) = \sum_{\mathbf{x} \in \mathbb{B}} \pi(\mathbf{x}) \kappa(\gamma \mid \mathbf{x})$
 $s \leftarrow 1$
repeat
 | **sample** $\mathbf{x}_k^{(s)} \sim \kappa(\bullet \mid \mathbf{x}_k^{(s-1)})$ **for all** $k \in \llbracket 1, n \rrbracket$
until $\zeta(\mathbf{X}^{(s)}) - \zeta(\mathbf{X}^{(s-1)}) < \zeta_\Delta^*$ **or** $\zeta(\mathbf{X}^{(s)}) > \zeta^*$
return $\mathbf{X}^{(s)} = (\mathbf{x}_1^{(s)}, \dots, \mathbf{x}_n^{(s)})^\top$

2.2.4 Transition kernels

Most transition kernels in Monte Carlo simulations are some variant of the Metropolis-Hastings kernel (see e.g. [Robert and Casella \(2004\)](#)),

$$\kappa_{t+1}(\gamma \mid \mathbf{x}) := \lambda_{q_{t+1}}(\gamma, \mathbf{x}) q_{t+1}(\gamma \mid \mathbf{x}) + \delta_{\mathbf{x}}(\gamma) \left[1 - \sum_{\mathbf{y} \in \mathbb{B}^d} \lambda_{q_{t+1}}(\mathbf{y}, \mathbf{x}) q_{t+1}(\mathbf{y} \mid \mathbf{x}) \right],$$

where we sample from the kernel by proposing a new state $\gamma \sim q_{t+1}(\gamma \mid \mathbf{x})$ and accepting the proposal with probability

$$\lambda_{q_{t+1}}(\gamma, \mathbf{x}) := 1 \wedge \frac{\tilde{\pi}_{t+1}(\gamma) q_{t+1}(\mathbf{x} \mid \gamma)}{\tilde{\pi}_{t+1}(\mathbf{x}) q_{t+1}(\gamma \mid \mathbf{x})} \quad (7)$$

or returning \mathbf{x} otherwise. Again, we denote by $\tilde{\pi}_{t+1} \propto \pi_{t+1}$ the unnormalized version of π_{t+1} since the kernel only requires the ratio of the unnormalized probability mass functions.

Symmetric kernel On binary spaces, a common choice for the proposal distribution is

$$q(\gamma \mid \mathbf{x}) = \sum_{k=1}^d p_k \delta_k(|\mathbf{x} - \gamma|) \frac{k!(d-k)!}{d!}, \quad (8)$$

with weight vector $\mathbf{p} \in [0, 1]^d$ normalized such that $|\mathbf{p}| = 1$.

With probability p_k , the kernel proposes a uniform draw from the k -neighborhood of \mathbf{x} ,

$$N_k(\mathbf{x}) := \{\gamma \in \mathbb{B}^d : |\mathbf{x} - \gamma| = k\}. \quad (9)$$

We refer to this type of kernel as *symmetric kernel* since $q(\gamma \mid \mathbf{x}) = q(\mathbf{x} \mid \gamma)$ and equation (7) simplifies. This class of kernels provide a higher mutation rate than the random-scan Gibbs kernel (see Schäfer and Chopin (2011) for a discussion).

Locally operating transition kernels of the symmetric type are known to be slowly mixing. If we put most weight on small values of k , the kernel only changes one or a few entries in each step. If we put more weight on larger values of k , the proposals will hardly ever be accepted if the invariant distribution π is multi-modal. Ideally, we want the particles sampled from the transition kernel to be nearly independent after a few move steps which is often hard to achieve using local transition kernels.

Adaptive independent kernel For the sequential Monte Carlo algorithm, we use *adaptive independent kernels* which have proposal distributions of the kind

$$q(\gamma \mid \mathbf{x}) = q_\theta(\gamma), \quad \theta \in \Theta,$$

which do not depend on the current state \mathbf{x} but have a parameter θ which we adapt during the course of the algorithm.

The adaptive independent kernel is rapidly mixing if we can fit the *parametric family* q_θ such that the proposal distribution $q_{t+1} = q_{\theta_{t+1}}$ is sufficiently close to the target distribution π_{t+1} , yielding thus, on average, high acceptance rates $\lambda_{q_{t+1}}$. The general idea behind this approach is to take the information gathered in the current particle approximation into account (see e.g. Chopin (2002)). The usefulness of this strategy for sampling on binary spaces has been illustrated by Schäfer and Chopin (2011).

We fit a parameter θ_{t+1} to the particle approximation of π_{t+1} according to some suitable criterion. Precisely, θ_{t+1} is taken to be the maximum likelihood or method of moments estimator applied to the weighted sample $(\mathbf{w}_{t+1}, \mathbf{X}_t)$. The choice of the parametric family q_θ is crucial to the implementation of a sequential Monte Carlo sampler with adaptive independent kernel. We discuss this issue in detail in Section 3.

Adaptation could, to a certain extent, also be done for local transition kernels. Nott and Kohn (2005) propose an adaptive kernel which replaces the full conditional distribution of the Gibbs sampler by an easy to compute linear approximation which is estimated from the sampled particles. This method accelerates Gibbs sampling if the target distribution π is hard to evaluate but does not provide fast mixing like the adaptive independent kernel (see Schäfer and Chopin (2011) for a comparison).

Still, the use of local kernels in the context of the proposed sequential Monte Carlo algorithm might be favorable if, for instance, the structure of the problem allows to rapidly compute the acceptance probabilities of local moves. Further, batches of local moves can be alternated with independent proposals to ensure that the algorithm explores the neighborhood of local modes sufficiently well.

2.3 Remark on discrete state spaces

Since the sample space \mathbb{B}^d is discrete, a given particle is not necessarily unique. This raises the question whether it is sensible to store multiple copies of the same weighted particle in our system. In the sequel, we discuss some more details concerning this issue which has only been touched upon briefly by Schäfer and Chopin (2011).

Let $n(\mathbf{x})$ denote the number of copies of the particle \mathbf{x} in the system (\mathbf{w}, \mathbf{X}) . Indeed, for parsimonious reasons, we could just keep a single representative of \mathbf{x} and aggregate the associated weights to $w_*(\mathbf{x}) = n(\mathbf{x}) w(\mathbf{x})$.

2.3.1 Impact on the effective sample size

Shifting weights between identical particles does not affect the nature of the particle approximation but it obviously changes the effective sample size $\eta_n(\mathbf{w})$ which is undesirable since we introduced the effective sample size as a criterion to measure the goodness of a particle approximation. From an aggregated particle system, we cannot distinguish the weight disparity induced by reweighting according to the importance function (4) and the weight disparity induced by multiple sampling of the same states which occurs if the mass of the target distribution is concentrated. More precisely, we cannot tell whether the effective sample size is actually due to the gap between π_t and π_{t+1} or the presence of particle copies due to the mass of π_t concentrating on a small proportion of the state space which occurs by construction of the auxiliary distribution in Section 1.2.

2.3.2 Impact on the resample-move step

Aggregating the weights means that the number of particles is not fixed at runtime. In this case, the straightforward way to implement the move step presented in Section 2.2.2 is breaking up the particles into multiple copies corresponding to their weights and moving them separately. But instead of permanently splitting and pooling the weights it seems more efficient to just keep the multiple copies.

We could, however, design a different kind of resample-move algorithm which first augments the number of particles in the move step and then resamples exactly n weighted particles from this extended system using a variant of the resampling procedure proposed by Fearnhead and Clifford (2003). A simple way to augment the number of particles is sampling and reweighting via

$$\mathbf{x}_k^{(1)} \sim q_{t+1}(\bullet \mid \mathbf{x}_k^{(0)}), \quad w_k^{(1)} = w_k \lambda, \quad w_k^{(0)} = w_k(1 - \lambda),$$

where $\lambda = \lambda_{q_{t+1}}(\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(0)})$ denotes the acceptance probability (7) of the Metropolis-Hastings kernel. We tested this variant but could not see any advantage over the standard sampler presented in the preceding sections. For the augment-resample type algorithm the implementation is more involved and the computational burden significantly higher. In particular, the Rao-Blackwellization effect one might achieve when replacing the accept-reject steps of the transition kernel by a single resampling step does not seem to justify the extra computational effort.

Indeed, aggregating the weights does not only prevent us from using the effective sample size criterion, but also requires extra computational time of $\mathcal{O}(n \log n)$ in each iteration of the move step since pooling the weights is as complex as sorting. With our application in mind, however, computational time is more critical than memory, and we therefore recommend to refrain from aggregating the weights.

3 Parametric families on binary spaces

We review three parametric families on \mathbb{B}^d . In contrast to the similar discussion in Schäfer and Chopin (2011), we also consider a parametric family which cannot be used in sequential Monte Carlo samplers but in the context of the cross-entropy method. For more details on parametric families on binary spaces we refer to Schäfer (2012).

3.1 Suitable parametric families

We frame some properties making a parametric family suitable as proposal distribution in sequential Monte Carlo algorithms.

- (a) For reasons of parsimony, we prefer a family of distributions with at most $d(d+1)/2$ parameters like the multivariate normal.
- (b) Given a sample $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ from the target distribution π , we need to estimate θ^* in a reasonable amount of computational time.
- (c) We need to generate samples $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_m)^\top$ from the family q_θ . We need the rows of \mathbf{Y} to be independent.
- (d) For the sequential Monte Carlo algorithm, we need to evaluate $q_\theta(\mathbf{y})$ point-wise. However, the cross-entropy method still works without this requirement.
- (e) We want the calibrated family q_{θ^*} to reproduce e.g. the marginals and covariance structure of π to ensure that the parametric family q_{θ^*} is sufficiently close to π .

3.2 Product family

The simplest non-trivial distributions on \mathbb{B}^d are certainly those having independent components.

Product family For a vector $\mathbf{m} \in (0, 1)^d$ of marginal probabilities, we define the *product family*

$$q_{\mathbf{m}}^\square(\boldsymbol{\gamma}) := \prod_{i=1}^d m_i^{\gamma_i} (1 - m_i)^{1-\gamma_i}. \quad (10)$$

3.2.1 Properties

We check the requirement list from Section 3.1: (a) The product family is parsimonious with $\dim(\theta) = d$. (b) The maximum likelihood estimator $\hat{\mathbf{m}}$ is the weighted sample mean. (c) We can easily sample $\mathbf{y} \sim q_{\mathbf{m}}^\square$. (d) We can easily evaluate the mass function $q_{\mathbf{m}}^\square(\mathbf{y})$. (e) However, the product family does not reproduce any dependencies we might observe in (\mathbf{w}, \mathbf{X}) .

The last point is the crucial weakness which makes the product family impractical for particle optimization algorithms on strongly multi-modal problems. Consequently, the rest of this section deals with ideas on how to sample binary vectors with a given dependence structure. There are, to our knowledge, two major strategies to this end.

- (1) We construct a generalized linear model which permits to compute the conditional distributions. We apply the chain rule and write q_θ as

$$q_\theta(\boldsymbol{\gamma}) = q_\theta(\gamma_1) \prod_{i=2}^d q_\theta(\gamma_i \mid \gamma_{1:i-1}), \quad (11)$$

which allows to sample the entries of a random vector component-wise.

- (2) We sample from an auxiliary distribution φ_θ and map the samples into \mathbb{B}^d . We call

$$q_\theta(\boldsymbol{\gamma}) = \int_{\tau^{-1}(\boldsymbol{\gamma})} \varphi_\theta(\mathbf{v}) d\mathbf{v} \quad (12)$$

a copula family, although we refrain from working with explicit uniform marginals.

We first present a generalized linear model and then review a copula approach.

3.3 Logistic conditionals family

Even for rather simple non-linear models we usually cannot derive closed-form expressions for the marginal probabilities required for sampling according to (11). Therefore, we might directly construct a parametric family from its conditional probabilities.

Logistic conditionals family We define, for a lower triangular matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$, the *logistic conditionals family* as

$$q_{\mathbf{A}}^{\ell}(\gamma) := \prod_{i \in [1, d]} \ell \left(a_{ii} + \sum_{j=1}^{i-1} a_{ij} \gamma_j \right)^{\gamma_i} \left[1 - \ell \left(a_{ii} + \sum_{j=1}^{i-1} a_{ij} \gamma_j \right) \right]^{1-\gamma_i}$$

where $\ell: \mathbb{R} \rightarrow (0, 1)$, $\ell(x) = [1 + \exp(-x)]^{-1}$ is the logistic function. We readily identify the product family $q_{\mathbf{m}}^{\square}$ as the special case $\mathbf{A} = \text{diag} \ell^{-1}(\mathbf{m})$.

The virtue of the logistic conditionals family is that, by construction, we can sample a random vector component-wise while the full probability $q_{\mathbf{A}}^{\ell}(\mathbf{y})$ of the sample \mathbf{y} is computed as a by-product of Procedure 3. We refer to the Online Supplement for instructions on how to fit the parameter \mathbf{A} .

Procedure 3: Sampling via chain rule factorization

```

 $\mathbf{y} = (0, \dots, 0)$ ,  $p \leftarrow 1$ 
for  $i \in [1, d]$  do
     $r \leftarrow q_{\mathbf{A}}^{\ell}(y_i = 1 \mid \mathbf{y}_{1:i-1}) = \ell(a_{ii} + \sum_{j=1}^{i-1} a_{ij} y_j)$ 
     $u \sim \mathcal{U}[0, 1]$ 
    if  $u < r$  then  $y_i \leftarrow 1$ 
     $p \leftarrow \begin{cases} p \cdot r & \text{if } y_i = 1 \\ p \cdot (1 - r) & \text{if } y_i = 0 \end{cases}$ 
end
return  $\mathbf{y}$ ,  $p$ 

```

3.3.1 Properties

We check the requirement list from Section 3.1: (a) The logistic conditionals family is sufficiently parsimonious with $\dim(\theta) = d(d+1)/2$. (b) We can fit the parameter \mathbf{A} via likelihood maximization. The fitting is computationally intensive but feasible. (c) We can sample $\mathbf{y} \sim q_{\mathbf{A}}^{\ell}$ using the chain rule factorization (11). (d) We can exactly evaluate $q_{\mathbf{A}}^{\ell}(\mathbf{y})$. (e) The family $q_{\mathbf{A}}^{\ell}$ reproduces the dependency structure of the data \mathbf{X} although we cannot explicitly compute the marginal probabilities.

3.4 Gaussian copula family

Let φ_{θ} be a family of multivariate auxiliary distributions on \mathbb{X} and $\tau: \mathbb{X} \rightarrow \mathbb{B}^d$ a mapping into the binary space. We can sample from the copula family (12) by setting $\mathbf{x} = \tau(\mathbf{v})$ for a draw $\mathbf{v} \sim \varphi_{\theta}$ from the auxiliary distribution. Most multivariate parametric families with at most $d(d+1)/2$ parameters appear to either have a rather limited dependency range or they do not scale to higher dimensions Joe (1996). Therefore, the natural and seemingly only viable option for φ_{θ} is the multivariate normal distribution Emrich and Piedmonte (1991).

Gaussian copula family For a vector $\mathbf{a} \in \mathbb{R}^d$ and a correlation matrix $\mathbf{\Sigma} \in \mathbb{R}^{d \times d}$, we introduce the mapping

$$\tau_{\mathbf{a}}: \mathbb{R}^d \rightarrow \mathbb{B}^d, \quad \tau_{\mathbf{a}}(\mathbf{v}) := (\mathbb{1}_{(-\infty, a_1]}(v_1), \dots, \mathbb{1}_{(-\infty, a_d]}(v_d)),$$

and define the *Gaussian copula family* as

$$q_{\mathbf{a}, \mathbf{\Sigma}}^{gc}(\gamma) := (2\pi)^{-\frac{d}{2}} \det \mathbf{\Sigma}^{-\frac{1}{2}} \int_{\tau_{\mathbf{a}}^{-1}(\gamma)} \exp\left(-\frac{1}{2} \mathbf{v}^\top \mathbf{\Sigma}^{-1} \mathbf{v}\right) d\mathbf{v}.$$

For index sets $I \subseteq \llbracket 1, d \rrbracket$, the cross-moments

$$m_I = \sum_{\gamma \in \mathbb{B}^d} q_{\mathbf{a}, \mathbf{\Sigma}}^{gc}(\gamma) \prod_{i \in I} \gamma_i$$

are equal the cumulative distribution function of the multivariate normal with respect to the entries indexed by I (see [Schäfer \(2012\)](#) for a more detailed discussion). In particular, the first and second moments are

$$m_i = \Phi_1(a_i), \quad m_{ij} = \Phi_2(a_i, a_j; \sigma_{ij}), \quad i, j \in \llbracket 1, d \rrbracket,$$

where $\Phi_1(\cdot)$ and $\Phi_2(\cdot, \cdot; \sigma_{ij})$ denote the cumulative distribution functions of the univariate and bivariate normal distributions with zero mean, unit variance and correlation coefficient $\sigma_{ij} \in [-1, 1]$. We refer to the Online Supplement for instructions on how to fit the parameters \mathbf{a} and $\mathbf{\Sigma}$.

3.4.1 Properties

We check the requirement list from Section 3.1: (a) The Gaussian copula family is sufficiently parsimonious with $\dim(\theta) = d(d+1)/2$. (b) We can fit the parameters \mathbf{a} and $\mathbf{\Sigma}$ via method of moments. However, the parameter $\mathbf{\Sigma}$ is not always positive definite. (c) We can sample $\mathbf{y} \sim q_{\mathbf{a}, \mathbf{\Sigma}}^{gc}$ using $\mathbf{y} = \tau_{\mathbf{a}}(\mathbf{v})$ with $\mathbf{v} \sim \varphi_{\mathbf{\Sigma}}$. (d) We cannot easily evaluate $q_{\mathbf{a}, \mathbf{\Sigma}}^{gc}(\mathbf{y})$ since this requires computing high-dimensional integral expressions which is a computationally challenging problem in itself (see e.g. [Genz and Bretz \(2009\)](#)). The Gaussian copula family is therefore less useful for sequential Monte Carlo samplers but can be incorporated into the cross-entropy method reviewed in Section 4.2. (e) The family $q_{\mathbf{a}, \mathbf{\Sigma}}^{gc}$ reproduces the exact mean and, possibly scaled, correlation structure.

3.5 Toy example

We briefly discuss a toy example to illustrate the usefulness of the parametric families. For the quadratic function

$$f(\mathbf{x}) = \mathbf{x}^\top \mathbf{F} \mathbf{x}, \quad \mathbf{F} := \begin{pmatrix} 1 & 2 & 1 & 0 \\ 2 & 1 & -3 & -2 \\ 1 & -3 & 1 & 2 \\ 0 & -2 & 2 & -2 \end{pmatrix}, \quad (13)$$

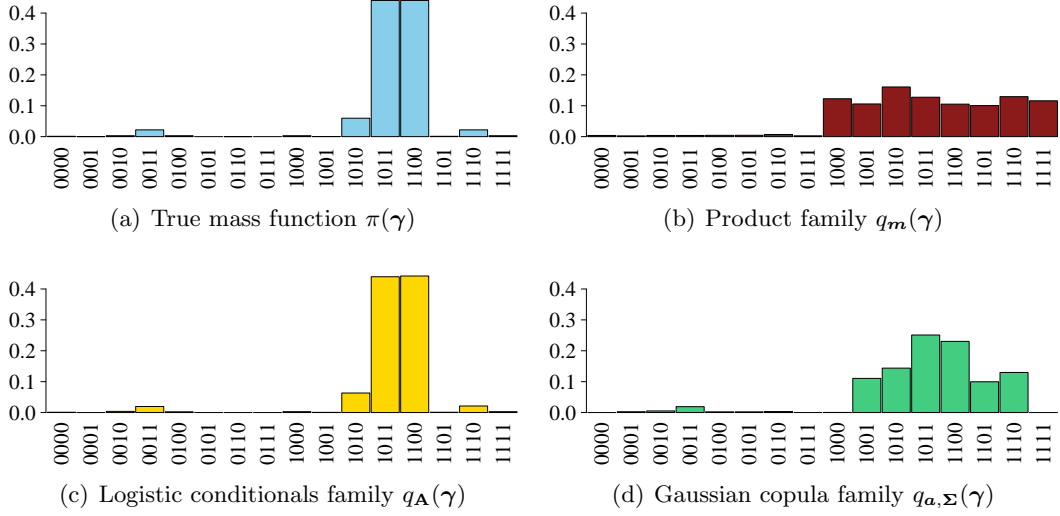
the associated probability mass function $\pi(\gamma) \propto \exp(\gamma^\top \mathbf{F} \gamma)$ has a correlation matrix

$$\mathbf{R} \approx \begin{pmatrix} 1 & 0.127 & -0.106 & -0.101 \\ 0.127 & 1 & -0.941 & -0.866 \\ -0.106 & -0.941 & 1 & 0.84 \\ -0.101 & -0.866 & 0.84 & 1 \end{pmatrix},$$

which indicates that this distribution has considerable dependencies and its mass function is therefore strongly multi-modal. We generate pseudo-random data from π , adjust the parametric families to the data and plot the mass functions of the fitted parametric families.

Figure 2 shows how the three parametric families cope with reproducing the true mass function. Clearly, the product family is not close enough to the true mass function to yield a suitable instrumental distribution while the logistic conditional family almost copies the characteristics of π and the Gaussian copula family allows for an intermediate goodness of fit.

Figure 2: Toy example showing how well the parametric families replicate the mass function of the distribution $\pi(\gamma) \propto \exp(\gamma^T \mathbf{F} \gamma)$ as defined in (13).



4 Optimization algorithms

In this section, we provide a synopsis of all steps involved in the sequential Monte Carlo algorithm and connect this framework to the cross-entropy method and simulated annealing. In Table 1, we state the necessary formulas for the tempered and the level set sequence introduced in Section 1.2.

4.1 Sequential Monte Carlo

For convenience, we summarize the complete sequential Monte Carlo sampler in Algorithm 1. Note that, in practice, the sequence π_{ϱ_t} is not indexed by t but rather by ϱ_t , which means that the counter t is only given implicitly.

The algorithm terminates if the particle diversity sharply drops below some threshold $\delta > 0$ which indicates that the mass has concentrated in a single mode. If we use a kernel with proposals from a parametric family q_{θ_t} , we might already stop if the family degenerates in the sense that only a few components of q_{θ_t} , say less than $d^* = 12$,

are random while the others are constant ones or zeros. In this situation, additional moves using a parametric family are a pointless effort. We either return the maximizer within the particle system or we solve the subproblem of dimension d^* by brute force enumeration. We might also perform some final local moves in order to further explore the regions of the state space the particles concentrated on.

Algorithm 1: Sequential Monte Carlo optimization

Input: $f: \mathbb{B}^d \rightarrow \mathbb{R}$
sample $\mathbf{x}_k \stackrel{\text{iid}}{\sim} \mathcal{U}_{\mathbb{B}^d}$ **for all** $k \in \llbracket 1, n \rrbracket$.
 $\alpha \leftarrow \text{find step length}(0, \mathbf{X})$,
 $\mathbf{w} \leftarrow \text{importance weights}(\alpha, \pi, \mathbf{X})$
while $\zeta_n(\mathbf{X}) > \delta$ **do**
 $q_\theta \leftarrow \text{fit parametric family}(\mathbf{w}, \mathbf{X})$ (see Section 3)
 $\hat{\mathbf{X}} \leftarrow \text{resample}(\mathbf{w}, \mathbf{X})$ (Procedure 1)
 $\mathbf{X} \leftarrow \text{move}(\kappa_{\pi, q_\theta}, \hat{\mathbf{X}})$ (Procedure 2)
 $\alpha \leftarrow \text{find step length}(\varrho, \mathbf{X})$
 $\mathbf{w} \leftarrow \text{importance weights}(\alpha, \pi_\varrho, \mathbf{X})$
 $\varrho \leftarrow \varrho + \alpha$
end
return $\arg\max_{\mathbf{x} \in \{\mathbf{x}_1, \dots, \mathbf{x}_n\}} f(\mathbf{x})$

4.2 Cross-entropy method

For the level set sequence, the effective sample size is the fraction of the particles which have an objective function value greater than $\max_{\mathbf{x} \in \mathbb{B}^d} f(\mathbf{x}) - 1/(\varrho_t + \alpha)$; see Table 1 and equation (3). The remaining particles are discarded since their weights equal zero. Consequently, there is no need to explicitly compute α_t as a solution of (6). We simply order the particles \mathbf{x}_k according to their objective values $f(\mathbf{x}_k)$ and only keep the $n(1-\beta)$ particles with the highest objective values.

Rubinstein [Rubinstein \(1997\)](#), who popularizes the use of level set sequences in the context of the cross-entropy method, refers to $n(1-\beta)$ as the size of the *elite sample*. The cross-entropy method has been applied successfully to a variety of combinatorial optimization problems, some of which are equivalent to pseudo-Boolean optimization [Rubinstein and Kroese \(2004\)](#), and is closely related to the proposed sequential Monte Carlo framework.

However, the central difference between the cross-entropy method and the sequential Monte Carlo algorithm outlined above is the use of the invariant transition kernel in the latter. We obtain the cross-entropy method as a special case if we replace the kernel κ_t by its proposal distribution q_{θ_t} . The sequential Monte Carlo approach uses a smooth family of distributions $\{\pi_\varrho: \varrho \geq 0\}$ and explicitly schedules the evolution π_{ϱ_t} which in turn leads to the proposal distributions q_{θ_t} . The cross-entropy method, in contrast, defines the subsequent proposal distribution

$$q_{\theta_{t+1}} \approx q_{\theta_t} \mathbb{1}_{L_{\varrho_{t+1}}^+}$$

without any reference sequence π_t to balance the speed of the particle evolution.

Table 1: Formulas for optimization sequences

	$\exp(\varrho f)$	$\mathbb{1}_{L_{\varrho}^+}$
$u_{t,\alpha}(\mathbf{x}_{k,t})$	$e^{\alpha f(\mathbf{x}_{k,t})}$	$\mathbb{1}_{L_{\varrho_t+\alpha}^+}(\mathbf{x}_{k,t})$
$\eta_n(\mathbf{w}_{t,\alpha})$	$\frac{[\sum_{k=1}^n e^{\alpha f(\mathbf{x}_{k,t})}]^2}{n \sum_{k=1}^n e^{2\alpha f(\mathbf{x}_{k,t})}}$	$\frac{ \{\mathbf{x}_{k,t} \mid k \in \llbracket 1, n \rrbracket\} \cap L_{\varrho_t+\alpha}^+ }{n}$
$\lambda_{q_{t+1}}(\gamma \mid \mathbf{x}_{k,t})$	$1 \wedge \frac{e^{\alpha(f(\gamma)-f(\mathbf{x}_{k,t}))}}{e^{\log q_t(\gamma)-\log q_t(\mathbf{x}_{k,t})}}$	$1 \wedge \frac{\mathbb{1}_{L_{\varrho_{t+1}}^+}(\gamma)}{e^{\log q_t(\gamma)-\log q_t(\mathbf{x}_{k,t})}}$

In order to decelerate the advancement of the cross-entropy method, we introduce a lag parameter $\tau \in [0, 1)$ and use a convex combination of the previous parameter θ_{t-1} and the parameter $\hat{\theta}_t$ fit to the current particle system, setting

$$\theta_t := (1 - \tau)\hat{\theta}_t + \tau\theta_{t-1}.$$

However, there are no guidelines on how to adjust the lag parameter during the run of the algorithm. Therefore, the sequential Monte Carlo algorithm is easier to calibrate since the reference sequence π_t controls the stride and automatically prevents the system from overshooting.

On the upside, the cross-entropy method allows for a broader class of auxiliary distributions q_{θ_t} since we do not need to evaluate $q_{\theta_t}(\mathbf{x})$ point-wise which is necessary in the computation of the acceptance probability of the Hastings kernel; see Section 3.4.

4.3 Simulated annealing

A well-studied approach to pseudo-Boolean optimization is simulated annealing [Kirkpatrick et al. \(1983\)](#). While the name stems from the analogy to the annealing process in metallurgy, there is a pure statistical meaning to this setup. We can picture simulated annealing as approximating the mode of a tempered sequence (2) using a single particle. Since a single observation does not allow for fitting a parametric family, we have to rely on symmetric transition kernels (8) in the move step.

A crucial choice is the sequence ϱ_t which in this context is often referred to as the *cooling schedule*. There is a vast literature advising on how to calibrate ϱ_t where a typical guideline is the expected acceptance rate of the Hastings kernel. We calibrate ϱ_t such that the empirical acceptance rate

$$\bar{\lambda}_{t-s:t} := \sum_{r=t-s}^t \lambda_r, \quad s > 0$$

follows approximately $(t+1)^{-5}$ for $t \in [0, 1]$. There are variants of simulated annealing which use more complex cooling schedules, tabu lists and multiple restarts, but we stick to this simple version for the sake of simplicity. Algorithm 2 describes the version we use in our numerical experiments in Section 5.3.3.

Algorithm 2: Simulated annealing optimization

```
[ht] Input:  $f: \mathbb{B}^d \rightarrow \mathbb{R}$ ,  $T^* \in \mathbb{R}$   
 $\mathbf{x} \sim \mathcal{U}_{\mathbb{B}^d}$ ,  $\mathbf{x}^* \leftarrow \mathbf{x}$ ,  $t \leftarrow 0$ ,  $T_\Delta \leftarrow 0$  (time elapsed)  
while  $T_\Delta < T^*$  do  
  sample  
     $\gamma \sim \mathcal{U}_{N_1(\mathbf{x})}$ ,  $u \sim \mathcal{U}_{[0,1]}$ ,  $\lambda_t \leftarrow 1 \wedge \exp[\varrho_t(f(\gamma) - f(\mathbf{x}))]$   
    if  $u < \lambda_t$  then  $\mathbf{x} \leftarrow \gamma$   
    if  $f(\mathbf{x}) > f(\mathbf{x}^*)$  then  $\mathbf{x}^* \leftarrow \mathbf{x}$   
    adjust  $\varrho_t$  such that  $\bar{\lambda}_{t-s:t} \approx (1 + T_\Delta/T^*)^{-5}$   
     $t \leftarrow t + 1$   
end  
return  $\mathbf{x}^*$ 
```

4.4 Randomized local search

We describe a greedy local search algorithm which works on any state space that allows for defining a neighborhood structure. The typical neighborhood on binary spaces is the k -neighborhood defined in (9). A greedy local search algorithm computes the objective value of all states in the current neighborhood and moves to the best state found until a local optimum is reached. The local search algorithm is called k -opt if it searches the neighborhood $\cup_{i=1}^k N_i(\cdot)$ (see e.g. Merz and Freisleben (2002) for a discussion).

The algorithm can be randomized by repeatedly restarting the procedure from randomly drawn starting points. There are more sophisticated versions of local search algorithms exploit the properties of the objective function but even a simple local search procedure can produce good results Alidaee et al. (2010). Algorithm 3 describes the 1-opt local search procedure we use in our numerical experiments in Section 5.3.3.

Algorithm 3: Randomized local search

```
[ht] Input:  $f: \mathbb{B}^d \rightarrow \mathbb{R}$ ,  $T^* \in \mathbb{R}$   
 $\mathbf{x}^* \sim \mathcal{U}_{\mathbb{B}^d}$ ,  $T_\Delta \leftarrow 0$  (time elapsed)  
while  $T_\Delta < T^*$  do  
   $\mathbf{x} \sim \mathcal{U}_{\mathbb{B}^d}$   
  while  $\mathbf{x}$  is not a local optimum do  
     $\mathbf{x} \leftarrow \operatorname{argmax}_{\gamma \in N_1(\mathbf{x})} f(\gamma)$   
  end  
  if  $f(\mathbf{x}) > f(\mathbf{x}^*)$  then  $\mathbf{x}^* \leftarrow \mathbf{x}$   
end  
return  $\mathbf{x}^*$ 
```

5 Applications

5.1 Unconstrained Quadratic Binary Optimization

5.1.1 Introduction

It is well-known that any pseudo-Boolean function $f: \mathbb{B}^d \rightarrow \mathbb{R}$ can be written as a multi-linear function

$$f(\mathbf{x}) = \sum_{I \subseteq \llbracket 1, d \rrbracket} f(\mathbf{1}_I(1), \dots, \mathbf{1}_I(d)) \prod_{i \in I} x_i \prod_{i \in \llbracket 1, d \rrbracket \setminus I} (1 - x_i) = \sum_{I \subseteq \llbracket 1, d \rrbracket} a_I \prod_{i \in I} x_i, \quad (14)$$

where $a_I \in \mathbb{R}$ are real-valued coefficients. We say the function f is of order k if the coefficients a_I are zero for all $I \subseteq \llbracket 1, d \rrbracket$ with $|I| > k$. While optimizing a first order function is trivial, optimizing a non-convex second order function is already an NP-hard problem [Garey and Johnson \(1979\)](#).

In the sequel, we focus on optimization of second order pseudo-Boolean functions to exemplify the stochastic optimization schemes discussed in the preceding sections. If f is a second order function, we restate program (1) as

$$\begin{aligned} & \text{maximize} && \mathbf{x}^\top \mathbf{F} \mathbf{x} \\ & \text{subject to} && \mathbf{x} \in \mathbb{B}^d, \end{aligned} \quad (15)$$

where $\mathbf{F} \in \mathbb{R}^{d \times d}$ is a symmetric matrix. We call (15) an unconstrained quadratic binary optimization problem (UQBO); we refer to [Boros et al. \(2007\)](#) for a list of applications and equivalent problems. In the literature it is also referred to as unconstrained quadratic Boolean or bivalent or zero-one programming [Beasley \(1998\)](#).

5.1.2 Particle optimization and meta-heuristics

Meta-heuristics are a class of algorithms that optimize a problem by improving a set of candidate solutions without systematically enumerating the state space; typically they deliver solutions in polynomial time while an exact solution has exponential worst case running time. The outcome is neither guaranteed to be optimal nor deterministic since most meta-heuristics are randomized algorithms. We briefly discuss the connection to particle optimization against the backdrop of the unconstrained quadratic binary optimization problem where we roughly separate them into two classes: local search algorithms and particle-driven meta-heuristics.

Local search algorithms iteratively improve the current candidate solution through local search heuristics and judicious exploration of the current neighborhood; examples are local search [Boros et al. \(2007\)](#); [Merz and Freisleben \(2002\)](#), tabu search [Glover et al. \(1998\)](#); [Palubeckis \(2004\)](#), simulated annealing [Katayama and Narihisa \(2001\)](#). Particle driven meta-heuristics propagate a set of candidate solutions and improve it through recombination and local moves of the particles; examples are genetic algorithms [Merz and Freisleben \(1999\)](#), memetic algorithms [Merz and Katayama \(2004\)](#), scatter search [Amini et al. \(1999\)](#). For comparisons of these methods we refer to [Hasan et al. \(2000\)](#) or [Beasley \(1998\)](#).

The sequential Monte Carlo algorithm and the cross-entropy method are clearly in the latter class of particle-driven meta-heuristics. The idea behind sequential Monte Carlo is closely related to the intuition behind population (or swarm) optimization and genetic

(or evolutionary) algorithms. However, the mathematical framework used in sequential Monte Carlo allows for a general formulation of the statistical properties of the particle evolution while genetic algorithms are often problem-specific and empirically motivated.

5.1.3 Particle optimization and exact solvers

If we can explicitly derive the multi-linear representation (14) of the objective function, there are techniques to turn program (1) into a linear program. For the UQBO it reads

$$\begin{aligned}
& \text{maximize} && f(\mathbf{x}) = 2 \sum_{i=1}^d \sum_{j=1}^{i-1} f_{ij} x_{ij} + \sum_{i=1}^d f_{ii} x_{ii} \\
& \text{subject to} && \mathbf{x} \in \mathbb{B}^{d(d+1)/2} \\
& && \left. \begin{aligned} x_{ij} &\leq x_{ii} \\ x_{ij} &\leq x_{jj} \\ x_{ij} &\geq x_{ii} + x_{jj} - 1 \end{aligned} \right\} \text{ for all } i, j \in \llbracket 1, d \rrbracket.
\end{aligned} \tag{16}$$

Note, however, that there are more parsimonious linearization strategies than this straightforward approach [(Hansen and Meyer, 2009), (Gueye and Michelon, 2009)]. The transformed problem allows to access the tool box of linear integer programming which consist of branch-and-bound algorithms that are combined with rounding heuristics, various relaxations techniques and cutting plane methods [(Pardalos and Rodgers, 1990), (Palubeckis, 1995)].

Naturally, the question arises whether particle-driven meta-heuristics can be incorporated into exact solvers to improve branch-and-bound algorithms. Indeed, stochastic meta-heuristics deliver lower bounds for maximization problems, but particle-driven algorithms are computationally somewhat expensive for this purpose unless the objective function is strongly multi-modal and other heuristics fail to provide good results; see the discussion in Section 5.2.5.

However, the sequential Monte Carlo approach in combination with the level set sequence (3) might also be useful to determine a global branching strategy, since the algorithm provides an estimator for

$$\bar{\gamma}_c := |L_c^+|^{-1} \sum_{\gamma \in \mathbb{B}^d} \gamma \mathbb{1}_{L_c^+}(\gamma),$$

which is the average of the super-level set $L_c^+ := \{\mathbf{x} \in \mathbb{B}^d : f(\mathbf{x}) \geq c\}$. These estimates given for a sequence of levels c might provide branching strategies than are superior to local heuristics or branching rules based on fractional solutions. A further discussion of this topic is beyond the scope of this paper but it certainly merits consideration.

5.2 Construction of test problems

5.2.1 Introduction

The meta-heuristics we want to compare do not exploit the quadratic structure of the objective function and might therefore be applied to any binary optimization program. If the objective function can be written in multi-linear form like (15) there are efficient local search algorithms Boros et al. (2007); Merz and Freisleben (2002) which exploit special properties of the target function and easily beat particle methods in terms of computational time.

Therefore, the use of particle methods is particularly interesting if the objective function is expensive to compute or even a black box. The posterior distribution in Bayesian variable selection for linear normal models is an example of such an objective function (see [Schäfer and Chopin \(2011\)](#) and references therein). We stick to the UQBO for our numerical comparison since problem instances of varying difficulty are easy to generate and interpret while the results carry over to general binary optimization.

In the vast literature on UQBO, authors typically compare the performance of meta-heuristics on a suite of randomly generated problems with certain properties. [Pardalos \(1991\)](#) proposes standardized performance tests on symmetric matrices $\mathbf{F} \in \mathbb{Z}^{d \times d}$ with entries f_{ij} drawn from the uniform

$$q_c(k) := \frac{1}{2c} \mathbb{1}_{[-c, c]}(k), \quad c \in \mathbb{N}.$$

The test suites generated by [Beasley \(1990, OR-library\)](#) and [Glover et al. \(1998\)](#) follow this approach have been widely used as benchmark problems in the UQBO literature (see [Boros et al. \(2007\)](#) for an overview). In the sequel we discuss the impact of diagonal dominance, shifts, the density and extreme values of \mathbf{F} on the expected difficulty of the corresponding UQBO problem.

5.2.2 Diagonal

Generally, stronger *diagonal dominance* in \mathbf{F} corresponds to easier UQBO problems [Billionnet and Sutter \(1994\)](#). Consequently, the original problem generator presented by [Pardalos \(1991\)](#) is designed to draw the off-diagonal elements from a uniform on a different support $\llbracket -q, q \rrbracket$ with $q \in \mathbb{N}$.

In this context, we point out that the impact of diagonal dominance carries over to the statistical properties of the tempered distributions (2) we defined in the introductory Section 1.2. Indeed, stronger diagonal dominance in \mathbf{F} corresponds to exponential quadratic distributions

$$\pi(\gamma) := \frac{\exp(\gamma^\top \mathbf{F} \gamma)}{\sum_{\gamma \in \mathbb{B}^d} \exp(\gamma^\top \mathbf{F} \gamma)}$$

having lower dependencies between the components of γ . We can analytically derive a parameter $\mathbf{A} \in \mathbb{R}^{d \times d}$ for a logistic conditionals family $q_{\mathbf{A}}^\ell$ that approximates $\pi(\gamma)$ where the quality of the approximation increases as the diagonal of \mathbf{F} becomes more dominant [Schäfer \(2012\)](#). We can accelerate the sequential Monte Carlo algorithm by initializing the system from $q_{\mathbf{A}}^\ell$ instead of $\mathcal{U}_{\mathbb{B}^d}$. However, we did not exploit this option to keep the present work more concise.

For positive definite $\mathbf{F} \succ 0$, the optimization problem is convex and can be solved in polynomial time [Kozlov et al. \(1979\)](#); in exact optimization, this fact is exploited to construct upper bounds for maximization problems [Poljak and Wolkowicz \(1995\)](#). We observe a corresponding complexity reduction in statistical modeling. For $\mathbf{F} \succ 0$, the auxiliary distribution

$$\pi(\gamma) := \frac{\gamma^\top \mathbf{F} \gamma}{2^{d-2} (\mathbf{1}^\top \mathbf{F} \mathbf{1} + \text{tr}(\mathbf{F}))},$$

is a feasible mass function, and we can derive analytical expressions concerning all cross-moments and marginal distributions [Schäfer \(2010\)](#) which allows to largely analyze the properties of $\pi(\gamma)$ without enumerating the state space.

5.2.3 Shifts

The global optimum of the UQBO problem is more difficult to detect as we shift the entries of the matrix \mathbf{F} but the relative gap between the optimum and any heuristic value diminishes. If we sample $f_{ij} = f_{ij}^\tau$ from a uniform on the *shifted* support

$$q_{c,\tau}(k) := \mathcal{U}_{[-c+\tau, c+\tau]}(k), \quad c \in \mathbb{N}, \tau \in \llbracket -c, c \rrbracket,$$

we obtain an objective function

$$f_\tau(\mathbf{x}) = \mathbf{x}^\top \mathbf{F}^\tau \mathbf{x} \stackrel{d}{=} \mathbf{x}^\top (\mathbf{F}^0 + \tau \mathbf{1}\mathbf{1}^\top) \mathbf{x} = f_0(\mathbf{x}) + \tau |\mathbf{x}|^2,$$

where $\stackrel{d}{=}$ means equality in distribution. Hence, with growing $|\tau|$ the optimum depends less on \mathbf{F} and the relative gap between the optimum and a solution provided by any meta-heuristic vanishes. Boros et al. (2007) define a related criterion

$$\bar{\rho} := \frac{1}{2} + \frac{\tau + 2\tau c}{2(\tau^2 + c^2 + c)} \in [0, 1]$$

and report a significant impact of $\bar{\rho}$ on the solution quality of their local search algorithms which is not surprising.

5.2.4 Density

The difficulty of the optimization problem is related to the number of interactions, that is the number of non-zero elements of \mathbf{F} . We call the proportion of non-zeros the *density* of \mathbf{F} . Drawing f_{ij} from the mixture

$$q_{c,\omega}(k) = \omega \mathcal{U}_{[-c,c]}(k) + (1 - \omega) \delta_0(k), \quad c \in \mathbb{N}, \omega \in (0, 1]$$

we adjust the difficulty of the problem to a given expected density ω .

Note that not all algorithms are equally sensitive to the density of \mathbf{F} . Using the basic linearization (16), each non-zero off-diagonal element requires the introduction of an auxiliary variable and three constraints. Thus, the expected total number of variables and the expected total number of constraints, which largely determine the complexity of the optimization problem, are proportional to the density ω .

On the other hand, many randomized approaches, including the particle algorithms discussed in Section 2, are less sensitive to the density of the problem in the sense that replacing zero elements by small values has a minor impact on the performance of these algorithms. Rather than the zero/non-zero duality, we suggest that the presence of extreme values determines the difficulty of providing heuristic solutions.

5.2.5 Extreme values

The uniform sampling approach advocated by Pardalos (1991) is widely used in the literature for comparing meta-heuristics. Certainly, particle-driven methods are computationally too expensive to outperform local search heuristics on test problems with uniformly drawn entries; Beasley (1998) confirms this intuition with respect to genetic algorithms versus tabu search and simulated annealing. However, the uniform distribution does not produce *extreme values* and it is vital to keep in mind that these have an enormous impact on the performance of local search algorithms.


Extreme values in \mathbf{F} lead to the existence of distinct local maxima $\mathbf{x}^* \in \mathbb{B}^d$ of f in the sense that there is no better candidate solution than \mathbf{x}^* in the neighborhood $\cup_{i=1}^k N_i(\mathbf{x}^*)$ even for relatively large k . Further, extreme local minima might completely prevent a local search heuristic from traversing the state space in certain directions. Consequently, local search algorithms, as discussed in Section 5.1.2, depend more heavily on their starting value, and their performance deteriorates with respect to particle-driven algorithms.

We propose to draw the matrix entries f_{ij} from a discretized Cauchy distribution

$$\mathcal{C}_c(k) \propto (1 + (k/c)^2)^{-1}, \quad c \in \mathbb{N} \quad (17)$$

that has heavy tails which cause extreme values to be frequently sampled. Figure 3 shows the distribution of a Cauchy and a uniform to illustrate the difference. The resulting UQBO problems have quite distinct local maxima; in that case we also say that the function $f(\mathbf{x})$ is *strongly multi-modal*.

Figure 3: Histograms of a Cauchy \mathcal{C}_5 and a uniform \mathcal{U}_{10} distribution.



5.3 Numerical comparison

In this section, we provide numerical comparisons based on instances of the UQBO problem. We generated two random test suites of dimension $d = 250$, each having 10 instances. For the first suite, we sampled the matrix entries from a uniform distribution \mathcal{U}_{100} on $\llbracket -100, 100 \rrbracket$; for the second, we sampled from a Cauchy distribution \mathcal{C}_{100} as defined in (17). For performance evaluation, we run a specified algorithm 100 times on the same problem and denote the outcome by $\mathbf{x}_1, \dots, \mathbf{x}_{100}$.

5.3.1 Visualization

Since the absolute values are not meaningful, we report the relative ratios

$$\varrho_k := \frac{f(\mathbf{x}_k) - \text{worst solution found}}{\text{best known solution} - \text{worst solution found}} \in [0, 1],$$

where the best known solution is the highest objective value ever found for that instance and the worst solution is the lowest objective value among the 100 outcomes. We summarize the results in a histogram. The first n bins are singletons $b_k := \{\varrho_k^*\}$ for the highest values $\varrho_1^* > \dots > \varrho_n^* \in \{\varrho_k : k \in \llbracket 1, 100 \rrbracket\}$; the following n bins are equidistant intervals $b_k^< := [\frac{n-k}{n}\varrho_n^*, \frac{n-k+1}{n}\varrho_n^*)$. The graphs show the bins $b_1, \dots, b_n, b_1^<, \dots, b_n^<$ in descending order from left to right on the x -axis. The interval bins are marked with a sign “<” and the lower bound. The y -axis represents the counts.

For comparison, we draw the outcome of several algorithms into the same histogram, where the worst solution found is the lowest overall objective value among the outcomes. For each algorithm, the counts are depicted in a different color and, for better readability,

with diagonal stripes in a different angle. To put it plainly, an algorithm performs well if its boxes are on the left of the graph since this implies that the outcomes were often close to the best known solution.


5.3.2 Comparison of binary parametric families

We study how the choice of the binary parametric family affects the quality of the delivered solutions. The focus is on the cross-entropy method, since we cannot easily use the Gaussian copula family in the context of sequential Monte Carlo. We use $n = 1.2 \times 10^4$ particles, set the speed parameter to $\beta = 0.8$ (or the elite fraction to 0.2) and the lag parameter to $\tau = 0.5$.

The numerical comparisons, given in Figures 5(b) and 5(a), clearly suggest that using more advanced binary parametric families allows the cross-entropy method to detect local maxima that are superior to those detected using the product family. Hence, the numerical experiments confirm the intuition of our toy example in Figure 2.

On the strongly multi-modal instance 5(a) the numerical evidence for this conjecture is stunningly clear-cut; on the weakly multi-modal problem 5(b) its validity is still unquestionable. This result seems natural since reproducing the dependencies induced by the objective function is more relevant in the former case than in the latter.

Figure 4: The cross-entropy method using different binary parametric families.



5.3.3 Comparison of optimization algorithms

We compare a sequential Monte Carlo sampler with parametric family, a sequential Monte Carlo sampler with a single-flip symmetric kernel (8), the cross-entropy method, simulated annealing and 1-opt local search as described in Section 4.

For the cross entropy method, we use the same parameters as in the preceding section. For the sequential Monte Carlo algorithm, we use $n = 0.8 \times 10^4$ particles and set the speed parameter to $\beta = 0.9$; we target a tempered auxiliary sequence (2). For both algorithms we use the logistic conditionals family as sampling distribution. With these configurations, the algorithms converge in roughly 25 minutes. We calibrate the sequential Monte Carlo sampler with local moves to have the same average run time by processing batches of 10 local moves before checking the particle diversity criterion. The simulated annealing and 1-opt local search algorithms run for exactly 25 minutes.

The results shown in Figures 6(b) and 6(a) assert the intuition that particle methods perform significantly better on strongly multi-modal problems. However, on the easy test problems, the particle methods tend to persistently converge to the same sub-optimal local modes. This effect is probably due to their poor local exploration properties. Since particle methods perform significantly less evaluations of the objective function, they are less likely to discover the highest peak in a region of rather flat local modes. The use of parametric families aggravates this effect, and it seems advisable to alternate global and local moves to make a particle algorithm more robust against this kind of behavior.

Further numerical results are shown in Figure 6 and Figure 7.

6 Discussion and conclusion


The numerical experiments carried out on different parametric families revealed that the use of the advanced families proposed in this paper significantly improves the performance of the particle algorithms, especially on the strongly multi-modal problems. The experiments demonstrate that local search algorithms, like simulated annealing and randomized 1-opt local search, indeed outperform particle methods on weakly multi-modal problems but deliver inferior results on strongly multi-modal problems.

Using tabu lists, adaptive restarts and rounding heuristics, we can certainly design local search algorithms that perform better than simulated annealing and 1-opt local search. Still, the structural problem of strong multi-modality persists for path-based algorithms. On the other hand, cleverly designed local search heuristics will clearly beat sequential Monte Carlo methods on easy to moderately difficult problems.

The results encourage the use of particle methods if the objective function is known to be potentially multi-modal and hard to analyze analytically. We have to keep in mind that multiple restarts of rather simple local search heuristics can be very efficient if they make use of the structure of the objective function. For 25 minutes of randomized restarts, the heuristic proposed by Boros et al. (2007), which exploits the fact that the partial derivatives of a multi-linear function are constant, practically always returns the best known solution on all test problems treated to create Figures 6 and 7.

The numerical work was completely done in Python 2.6 using SciPy packages and run on a cluster with 1.86 GHz processors. The sources used in this work and the problems processed in this paper can be found at <http://code.google.com/p/smcdss>.

Figure 5: Comparison of stochastic optimization algorithms on two UQBO problems.



7 Acknowledgments

This work is part of the author's Ph.D. thesis at CREST under supervision of Nicolas Chopin whom I would like to thank for the numerous discussions on particle algorithms. I thank the editor and two anonymous referees for their detailed comments which helped to significantly improve this paper.

8 Appendix: Fitting the parameters

We briefly summarize how the parameters of the logistic conditionals family and the Gaussian copula family can be assessed for a given particle system $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ with weights $\mathbf{w} = (w_1, \dots, w_n)$. We denote by

$$\bar{x}_i := \sum_{k=1}^n w_k x_{ki}, \quad \bar{x}_{ij} := \sum_{k=1}^n w_k x_{ki} x_{kj}, \quad i, j \in \llbracket 1, d \rrbracket \quad (18)$$

the weighted first and second sample moments.

Figure 6: Comparison of stochastic optimization algorithms. 10 problems $f(\mathbf{x}) = \mathbf{x}^\top \mathbf{F} \mathbf{x}$ with $f_{ij} \sim \mathcal{C}_{100}$ for $i, j \in \llbracket 1, 250 \rrbracket$

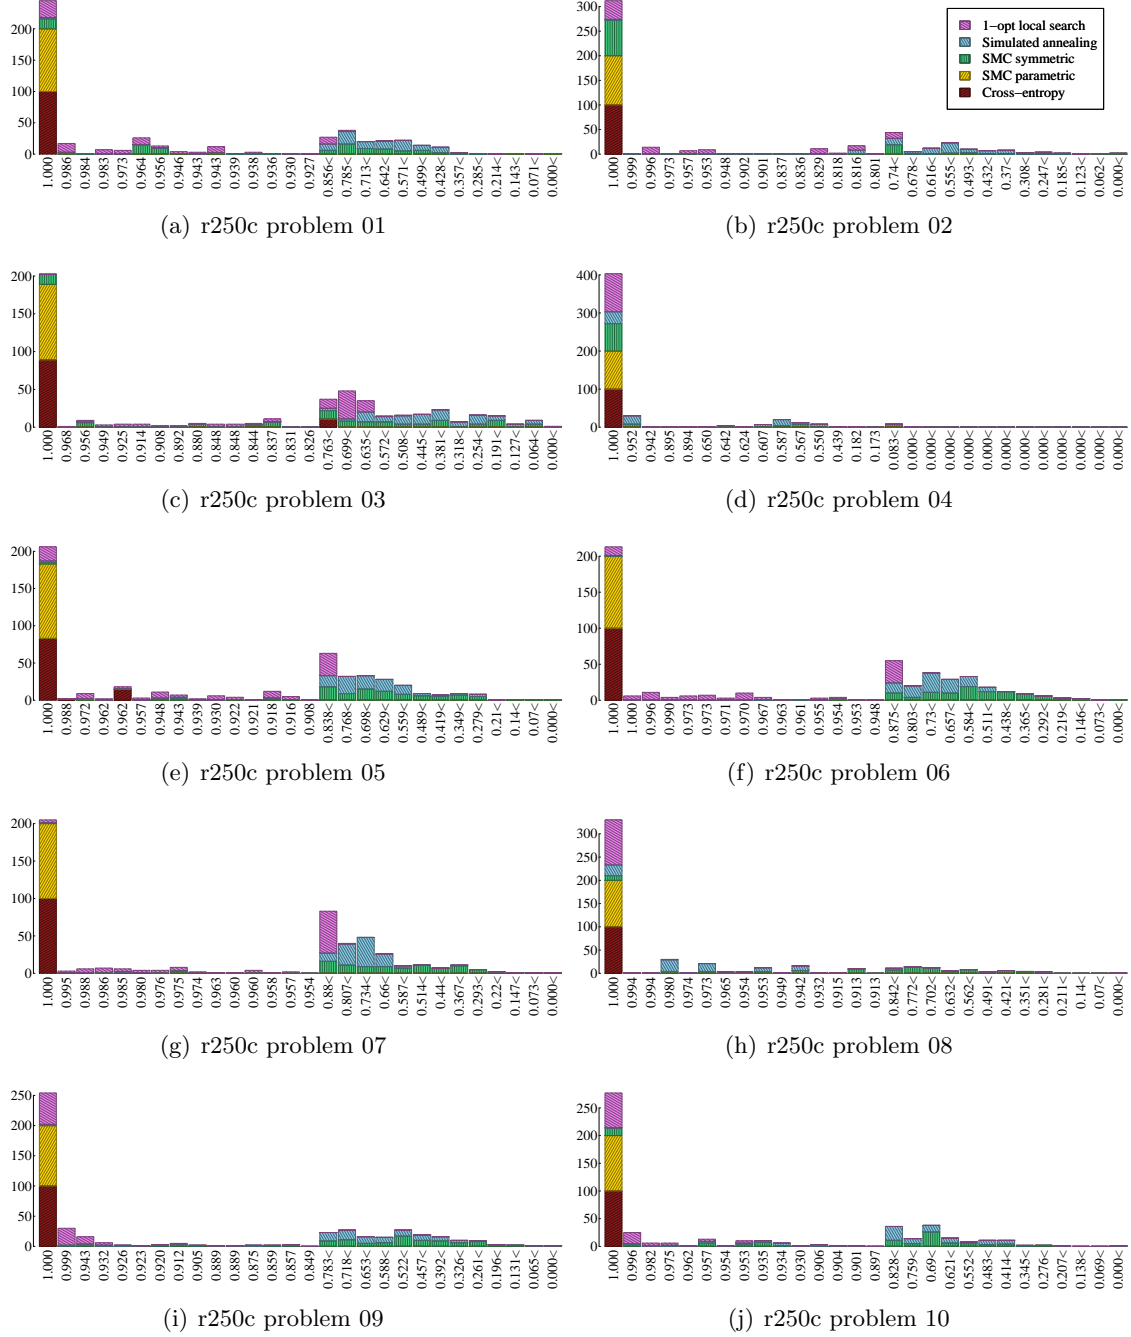
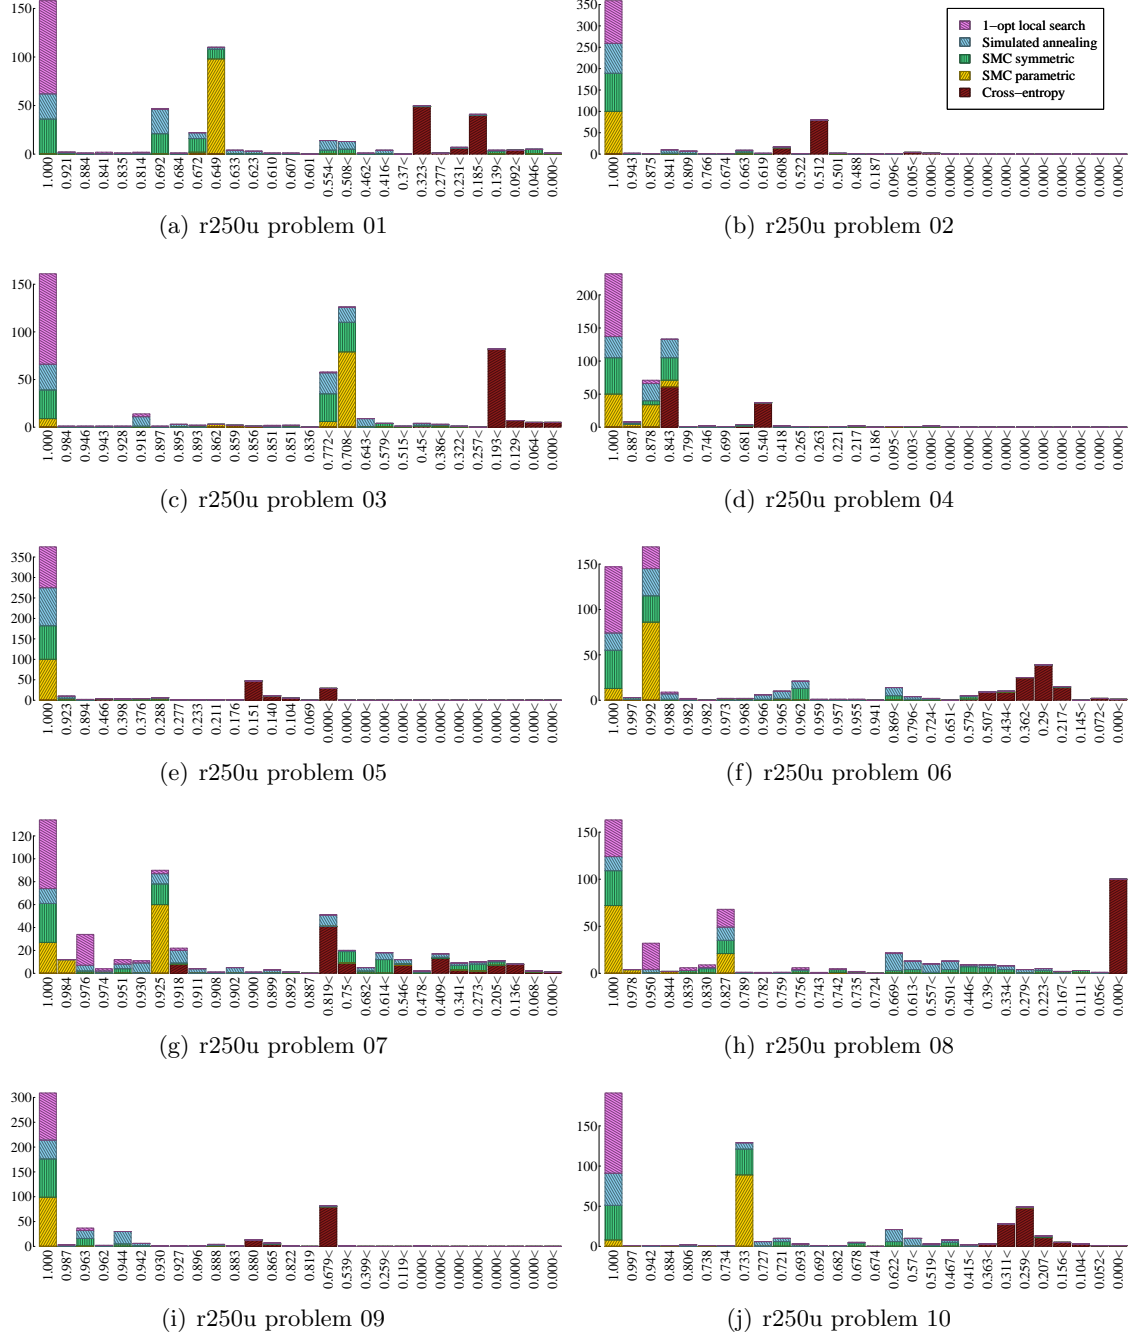


Figure 7: Comparison of stochastic optimization algorithms. 10 problems $f(\mathbf{x}) = \mathbf{x}^\top \mathbf{F} \mathbf{x}$ with $f_{ij} \sim \mathcal{U}_{100}$ for $i, j \in \llbracket 1, 250 \rrbracket$



8.1 Logistic conditionals family

8.1.1 Derivatives of the log-likelihood function

The log-likelihood function of the weighted logistic regression of $\mathbf{y}^{(i)} := \mathbf{X}_{\bullet i}$ on $\mathbf{Z}^{(i)} := (\mathbf{X}_{\bullet 1:i-1}, \mathbf{1})$ is

$$\begin{aligned} \log L(\mathbf{a}) &= \sum_{k=1}^n w_k \left[y_k^{(i)} \log[\ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a})] + (1 - y_k^{(i)}) \log[1 - \ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a})] \right] \\ &= \sum_{k=1}^n w_k \left[y_k^{(i)} \mathbf{z}_{k\bullet}^{(i)} \mathbf{a} - \log[1 + \exp(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a})] \right], \end{aligned}$$

where we used that $\log[1 - \ell(\mathbf{x}^\top \mathbf{a})] = -\log[1 + \exp(\mathbf{x}^\top \mathbf{a})] = -\mathbf{x}^\top \mathbf{a} + \log[\ell(\mathbf{x}^\top \mathbf{a})]$. Since $\partial \log[1 + \exp(\mathbf{x}^\top \mathbf{a})] / \partial \mathbf{a} = \ell(\mathbf{x}^\top \mathbf{a}) \mathbf{x}$, the gradient of the log-likelihood is

$$s(\mathbf{a}) = \sum_{k=1}^n w_k \left[y_k^{(i)} \mathbf{z}_{k\bullet}^{(i)} - \ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a}) \mathbf{z}_{k\bullet}^{(i)} \right] = (\mathbf{Z}^{(i)})^\top \text{diag}(\mathbf{w}) [\mathbf{y}^{(i)} - \mathbf{p}_\mathbf{a}^{(i)}],$$

where $(p_\mathbf{a}^{(i)})_k := \ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a})$. Since $\partial \ell(\mathbf{x}^\top \mathbf{a}) / \partial \mathbf{a} = -\ell(\mathbf{x}^\top \mathbf{a}) [1 - \ell(\mathbf{x}^\top \mathbf{a})] \mathbf{x}$, the Hessian matrix of the log-likelihood is

$$s'(\mathbf{a}) = - \sum_{k=1}^n w_k \left[\ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a}) [1 - \ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a})] \right] \mathbf{z}_{k\bullet}^{(i)} (\mathbf{z}_{k\bullet}^{(i)})^\top = -(\mathbf{Z}^{(i)})^\top \text{diag}(\mathbf{w}) \text{diag}(\mathbf{q}_\mathbf{a}^{(i)}) \mathbf{Z}^{(i)},$$

where $(q_\mathbf{a}^{(i)})_k := \ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a}) [1 - \ell(\mathbf{z}_{k\bullet}^{(i)} \mathbf{a})]$.

8.1.2 Complete separation

The data might suffer from complete or quasi-complete separation [Albert and Anderson \(1984\)](#) which causes the likelihood function $L(\mathbf{a})$ to be monotonic. In that case there is no maximizer. We can avoid the monotonicity by assigning a suitable prior distribution to the parameter \mathbf{a} . [Firth \(1993\)](#) recommends the Jeffreys prior for its bias reduction which can conveniently be implemented via data adjustment [Kosmidis and Firth \(2009\)](#).

For the sake of simplicity, however, we only assign a simple Gaussian prior with variance $\varepsilon^{-1} > 0$ such that, up to a constant, the log-posterior distribution is the log-likelihood function plus a quadratic penalty term and therefore always convex. The score function and its Jacobian matrix become

$$s(\mathbf{a}) = (\mathbf{Z}^{(i)})^\top \text{diag}(\mathbf{w}) [\mathbf{y}^{(i)} - \mathbf{p}_\mathbf{a}^{(i)}] - \varepsilon \mathbf{a}, \quad s'(\mathbf{a}) = -(\mathbf{Z}^{(i)})^\top \text{diag}(\mathbf{w}) \text{diag}(\mathbf{q}_\mathbf{a}^{(i)}) \mathbf{Z}^{(i)} - \varepsilon \mathbf{I}.$$

The bias-reduced estimators are known to shrink towards 0.5 which is an undesired property when fitting a parametric family. Therefore, we attempt to keep the shrinkage parameter ε as small as possible.

8.1.3 Newton-Raphson iterations

The first order condition $s(\mathbf{a}) = \mathbf{0}$ is solved iteratively

$$\mathbf{a}^{(t+1)} = \mathbf{a}^{(t)} - [s'(\mathbf{a}^{(t)})]^{-1} s(\mathbf{a}^{(t)}) = \mathbf{a}^{(t)} + \mathbf{x}^{(t)}$$

where $\varphi(\cdot, \cdot; \sigma)$ denotes the density of the bivariate normal distribution with correlation coefficient σ . We obtain the following Newton-Raphson iteration

$$\sigma^{(t+1)} = \sigma^{(t)} - \frac{\Phi_2(a_i, a_j; \sigma_{ij}^{(t)}) - \bar{x}_{ij}}{\varphi(a_i, a_j; \sigma_{ij}^{(t)})}, \quad (21)$$

starting at some initial value $\sigma_{ij}^{(0)} \in (-1, 1)$; see Procedure 5. In the sequential Monte Carlo context, good initial values are obtained from the parameters of the previous auxiliary distributions.

We use a fast series approximation [Drezner and Wesolowsky \(1990\)](#) to evaluate $\Phi_2(a_i, a_j; \sigma_{ij})$. These approximations are critical when σ_{ij} comes very close to either boundary of $[-1, 1]$. The Newton iteration might repeatedly fail when restarted at the corresponding boundary $\sigma_{ij}^{(0)} \in \{-1, 1\}$. In any event, $\Phi_2(x_1, x_2; \sigma)$ is strictly monotonic in σ since its derivative (20) is positive, and we can switch to bi-sectional search if necessary.

Algorithm 5: Fitting the dependency matrix

Input: \bar{x}_i, \bar{x}_{ij} **for all** $i, j \in \llbracket 1, d \rrbracket$
 $a_i \leftarrow \Phi^{-1}(\bar{x}_i)$ **for all** $i \in \llbracket 1, d \rrbracket$
 $\Sigma^{(0)} \leftarrow \mathbf{I}$
for $i, j \in \llbracket 1, d \rrbracket, i < j$ **do**
 repeat
 $\sigma_{ij}^{(t+1)} \leftarrow \sigma_{ij}^{(t)} - \frac{\Phi_2(a_i, a_j; \sigma_{ij}^{(t)}) - \bar{x}_{ij}}{\varphi(a_i, a_j; \sigma_{ij}^{(t)})}$
 until $|\sigma_{ij}^{(t+1)} - \sigma_{ij}^{(t)}| < \delta$
 $\sigma_{ji} \leftarrow \sigma_{ij}^{(t+1)}$
end
if not $\Sigma \succ 0$ **then** $\Sigma \leftarrow (\Sigma + |\lambda| \mathbf{I}) / (1 + |\lambda|)$
return \mathbf{a}, Σ

The locally fitted correlation matrices Σ might not be positive definite for $d \geq 3$ because the Gaussian copula is not flexible enough to model the full range of cross-moments binary distributions might have (see [Schäfer \(2012\)](#) for an extended discussion). We obtain a feasible parameter replacing Σ by

$$\Sigma^* = (\Sigma + |\lambda| \mathbf{I}) / (1 + |\lambda|),$$

where λ is smaller than all eigenvalues of the locally fitted matrix Σ . This approach evenly lowers the local correlations to a feasible level and is easy to implement on standard software. In practice, we also prefer to work with a sparse version of the Gaussian copula family that concentrates on strong dependencies and sets minor correlations to zero.

References

Albert, A. and Anderson, J. A. (1984). On the existence of maximum likelihood estimates in logistic regression models. *Biometrika*, 72:1–10.

- Alidaee, B., Kochenberger, G., and Wang, H. (2010). Theorems supporting r-flip search for pseudo-Boolean optimization. *International Journal of Applied Metaheuristic Computing (IJAMC)*, 1(1):93–109.
- Amini, M., Alidaee, B., and Kochenberger, G. (1999). A scatter search approach to unconstrained quadratic binary programs. In *New ideas in optimization*, pages 317–330. McGraw-Hill Ltd., UK.
- Beasley, J. (1990). OR-Library: Distributing test problems by electronic mail. *Journal of the Operational Research Society*, pages 1069–1072.
- Beasley, J. (1998). Heuristic algorithms for the unconstrained binary quadratic programming problem. Technical report, Management School, Imperial College London.
- Billionnet, A. and Sutter, A. (1994). Minimization of a quadratic pseudo-Boolean function. *European Journal of Operational Research*, 78(1):106–115.
- Boros, E. and Hammer, P. (2002). Pseudo-Boolean optimization. *Discrete Applied Mathematics*, 123(1-3):155–225.
- Boros, E., Hammer, P., and Tavares, G. (2007). Local search heuristics for quadratic unconstrained binary optimization (QUBO). *Journal of Heuristics*, 13(2):99–132.
- Carpenter, J., Clifford, P., and Fearnhead, P. (1999). Improved Particle Filter for non-linear problems. *IEE Proc. Radar, Sonar Navigation*, 146(1):2–7.
- Chopin, N. (2002). A sequential particle filter method for static models. *Biometrika*, 89(3):539.
- Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436.
- Del Moral, P., Doucet, A., and Jasra, A. (2011). An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statistics and Computing*, to appear. doi: 10.1007/s11222-011-9271-y.
- Drezner, Z. and Wesolowsky, G. O. (1990). On the computation of the bivariate normal integral. *Journal of Statistical Computation and Simulation*, 35:101–107.
- Emrich, L. and Piedmonte, M. (1991). A method for generating high-dimensional multivariate binary variates. *The American Statistician*, 45:302–304.
- Fearnhead, P. and Clifford, P. (2003). Online inference for hidden Markov models via particle filters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(4):887–899.
- Firth, D. (1993). Bias reduction of maximum likelihood estimates. *Biometrika*, 80:27–38.
- Garey, M. and Johnson, D. (1979). *Computers and Intractability: A Guide to the Theory of NP-completeness*. WH Freeman & Co.
- Genz, A. and Bretz, F. (2009). *Computation of multivariate normal and t probabilities*, volume 195. Springer.

- Glover, F., Kochenberger, G., and Alidaee, B. (1998). Adaptive memory tabu search for binary quadratic programs. *Management Science*, 44:336–345.
- Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. Radar, Sonar Navigation*, 140(2):107–113.
- Gueye, S. and Michelon, P. (2009). A linearization framework for unconstrained quadratic (0-1) problems. *Discrete Applied Mathematics*, 157(6):1255–1266.
- Hansen, P. and Meyer, C. (2009). Improved compact linearizations for the unconstrained quadratic 0-1 minimization problem. *Discrete Applied Mathematics*, 157(6):1267–1290.
- Hasan, M., AlKhamis, T., and Ali, J. (2000). A comparison between simulated annealing, genetic algorithm and tabu search methods for the unconstrained quadratic Pseudo-Boolean function. *Computers & Industrial Engineering*, 38(3):323–340.
- Jasra, A., Stephens, D., Doucet, A., and Tsagaris, T. (2011). Inference for Lévy-Driven Stochastic Volatility Models via Adaptive Sequential Monte Carlo. *Scandinavian Journal of Statistics*.
- Joe, H. (1996). Families of m-variate distributions with given margins and m (m-1)/2 bivariate dependence parameters. *Lecture Notes-Monograph Series*, 28:120–141.
- Johansen, A., Moral, P. D., and Doucet, A. (2006). Sequential Monte Carlo samplers for rare events. In *Proceedings of the 6th International Workshop on Rare Event Simulation*, pages pp. 256–267.
- Johnson, N., Kotz, S., and Balakrishnan, N. (2002). *Continuous multivariate distributions - models and applications*, volume 2. New York: John Wiley & Sons,.
- Katayama, K. and Narihisa, H. (2001). Performance of Simulated Annealing-based heuristic for the unconstrained binary quadratic programming problem. *E. J. of Operational Research*, 134(1):103–119.
- Kirkpatrick, S., Gelatt, C., and Vecchi, M. (1983). Optimization by simulated annealing. *Science*, 220(4598):671.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25.
- Kong, A., Liu, J. S., and Wong, W. H. (1994). Sequential imputation and Bayesian missing data problems. *Journal of the American Statistical Association*, 89:278–288.
- Kosmidis, I. and Firth, D. (2009). Bias reduction in exponential family nonlinear models. *Biometrika*, 96(4):793–804.
- Kozlov, M., Tarasov, S., and Khachiyan, L. (1979). Polynomial solvability of convex quadratic programming. In *Soviet Mathematics Doklady*, volume 20, pages 1108–1111.
- Liu, J. and Chen, R. (1998). Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93(443):1032–1044.

- Merz, P. and Freisleben, B. (1999). Genetic algorithms for binary quadratic programming. In *Proceedings of the genetic and evolutionary computation conference*, volume 1, pages 417–424. Citeseer.
- Merz, P. and Freisleben, B. (2002). Greedy and local search heuristics for unconstrained binary quadratic programming. *Journal of Heuristics*, 8(2):197–213.
- Merz, P. and Katayama, K. (2004). Memetic algorithms for the unconstrained binary quadratic programming problem. *BioSystems*, 78(1-3):99–118.
- Nott, D. and Kohn, R. (2005). Adaptive sampling for Bayesian variable selection. *Biometrika*, 92(4):747.
- Palubeckis, G. (1995). A heuristic-based branch and bound algorithm for unconstrained quadratic zero-one programming. *Computing*, 54(4):283–301.
- Palubeckis, G. (2004). Multistart tabu search strategies for the unconstrained binary quadratic optimization problem. *Annals of Operations Research*, 131(1):259–282.
- Pardalos, P. (1991). Construction of test problems in quadratic bivalent programming. *ACM Transactions on Mathematical Software (TOMS)*, 17(1):74–87.
- Pardalos, P. and Rodgers, G. (1990). Computational aspects of a branch and bound algorithm for quadratic zero-one programming. *Computing*, 45(2):131–144.
- Poljak, S. and Wolkowicz, H. (1995). Convex relaxations of $(0, 1)$ -quadratic programming. *Mathematics of Operations Research*, pages 550–561.
- Robert, C. and Casella, G. (2004). *Monte Carlo statistical methods*. Springer Verlag.
- Rubinstein, R. Y. (1997). Optimization of computer simulation models with rare events. *European Journal of Operations Research*, 99:89–112.
- Rubinstein, R. Y. and Kroese, D. P. (2004). *The Cross-Entropy Method: A unified approach to combinatorial optimization, Monte-Carlo simulation, and Machine Learning*. Springer-Verlag.
- Schäfer, C. (2010). Parametric families on binary spaces. Technical report. arXiv:1008.0055.
- Schäfer, C. (2012). On parametric families for sampling binary data with specified mean and correlation. Technical report. arXiv:1111.0576.
- Schäfer, C. and Chopin, N. (2011). Sequential Monte Carlo on large binary sampling spaces. *Statistics and Computing*, to appear. doi: 10.1007/s11222-011-9299-z.